

Distributed Methods for High-dimensional and Large-scale Tensor Factorization

Kijung Shin

Dept. of Computer Science and Engineering
Seoul National University
Seoul, Republic of Korea
koreaskj@snu.ac.kr

U Kang

Dept. of Computer Science
KAIST
Daejeon, Republic of Korea
ukang@cs.kaist.ac.kr

Abstract—Given a high-dimensional large-scale tensor, how can we decompose it into latent factors? Can we process it on commodity computers with limited memory? These questions are closely related to recommender systems, which have modeled rating data not as a matrix but as a tensor to utilize contextual information such as time and location. This increase in the dimension requires tensor factorization methods scalable with both the dimension and size of a tensor. In this paper, we propose two distributed tensor factorization methods, SALS and CDTF. Both methods are scalable with all aspects of data, and they show an interesting trade-off between convergence speed and memory requirements. SALS updates a subset of the columns of a factor matrix at a time, and CDTF, a special case of SALS, updates one column at a time. In our experiments, only our methods factorize a 5-dimensional tensor with 1 billion observable entries, 10M mode length, and 1K rank, while all other state-of-the-art methods fail. Moreover, our methods require several orders of magnitude less memory than our competitors. We implement our methods on MAPREDUCE with two widely-applicable optimization techniques: local disk caching and greedy row assignment. They speed up our methods up to 98.2 \times and also the competitors up to 5.9 \times .

Keywords—Tensor factorization; Recommender system; Distributed computing; MapReduce

I. INTRODUCTION AND RELATED WORK

Recommendation problems can be viewed as completing a partially observable user-item matrix whose entries are ratings. Matrix factorization (MF), which decomposes the input matrix into a user factor matrix and an item factor matrix such that their multiplication approximates the input matrix, is one of the most widely-used methods for matrix completion [1], [2], [3]. To handle web-scale data, efforts were made to find distributed ways for MF, including ALS [3], DSGD [4], and CCD++ [5].

On the other hand, there have been attempts to improve the accuracy of recommendation by using additional contextual information such as time and location. A straightforward way to utilize such extra factors is to model rating data as a partially observable tensor where additional dimensions correspond to the extra factors. Similar to the matrix completion, tensor factorization (TF), which decomposes the input tensor into multiple factor matrices and a core tensor, has been used for tensor completion [6], [7], [8].

Table I: Summary of scalability results. The factors which each method is scalable with are checked. Our proposed SALS and CDTF are the only methods scalable with all the factors.

	CDTF	SALS	ALS	PSGD	FLEXIFACT
Dimension	✓	✓	✓	✓	
Observations	✓	✓	✓	✓	✓
Mode Length	✓	✓			✓
Rank	✓	✓			✓
Machines	✓	✓	✓		

As the dimension of web-scale recommendation problems increases, a necessity for TF algorithms scalable with the dimension as well as the size of data has arisen. A promising way to find such algorithms is to extend distributed MF algorithms to higher dimensions. However, the scalability of existing methods including ALS [3], PSGD [9], and FLEXIFACT [10] is limited as summarized in Table I.

In this paper, we propose Subset Alternating Least Square (SALS) and Coordinate Descent for Tensor Factorization (CDTF), distributed tensor factorization methods scalable with all aspects of data. SALS updates a subset of the columns of a factor matrix at a time, and CDTF, a special case of SALS, updates one column at a time. These two methods have distinct advantages: SALS converges faster, and CDTF is more memory-efficient. Our methods can also be used in any applications handling large-scale partially observable tensors, including social network analysis [11] and Web search [12].

The main contributions of our study are as follows:

- **Algorithm.** We propose SALS and CDTF, scalable tensor factorization algorithms. Their distributed versions are the only methods scalable with all the following factors: the dimension and size of data, the number of parameters, and the number of machines (Table I).
- **Analysis.** We analyze our methods and the competitors in a general N-dimensional setting in the following aspects: computational complexity, communication complexity, memory requirements, and convergence speed (Table II).
- **Optimization.** We implement our methods on MAPREDUCE with two widely-applicable optimization techniques: local disk caching and greedy row assignment.

Table II: Summary of distributed tensor factorization algorithms. The performance bottlenecks which prevent each algorithm from handling web-scale data are colored red. Only our proposed SALS and CDTF have no bottleneck. Communication complexity is measured by the number of parameters that each machine exchanges with the others. For simplicity, we assume that workload of each algorithm is equally distributed across machines, that the length of every mode is equal to I , and that T_{in} of SALS and CDTF is set to one.

Algorithm	Computational complexity (per iteration)	Communication complexity (per iteration)	Memory requirements	Convergence speed
CDTF	$O(\Omega N^2K/M)$	$O(NIK)$	$O(NI)$	Fast
SALS	$O(\Omega NK(N+C)/M + NIKC^2/M)$	$O(NIK)$	$O(NIC)$	Fastest
ALS [3]	$O(\Omega NK(N+K)/M + NIK^3/M)$	$O(NIK)$	$O(NIK)$	Fastest
PSGD [9]	$O(\Omega NK/M)$	$O(NIK)$	$O(NIK)$	Slow
FLEXIFACT [10]	$O(\Omega NK/M)$	$O(M^{N-2}NIK)$	$O(NIK/M)$	Fast

Table III: Table of symbols.

Symbol	Definition
\mathcal{X}	input tensor ($\in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$)
$x_{i_1 \dots i_N}$	(i_1, \dots, i_N) th entry of \mathcal{X}
N	dimension of \mathcal{X}
I_n	length of the n th mode of \mathcal{X}
$\mathbf{A}^{(n)}$	n th factor matrix ($\in \mathbb{R}^{I_n \times K}$)
$a_{i_n k}^{(n)}$	(i_n, k) th entry of $\mathbf{A}^{(n)}$
K	rank of \mathcal{X}
Ω	set of indices of observable entries of \mathcal{X}
$\Omega_{i_n}^{(n)}$	subset of Ω whose n th mode's index is equal to i_n
mS_n	set of rows of $\mathbf{A}^{(n)}$ assigned to machine m
\mathcal{R}	residual tensor ($\in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N}$)
$r_{i_1 \dots i_N}$	(i_1, \dots, i_N) th entry of \mathcal{R}
M	number of machines (reducers on MAPREDUCE)
T_{out}	number of outer iterations
T_{in}	number of inner iterations
λ	regularization parameter
C	number of parameters updated at a time
η_0	initial learning rate

They speed up not only our methods (up to $98.2\times$) but also the competitors (up to $5.9\times$) (Figure 6).

- **Experiment.** We empirically confirm the superior scalability of our methods and their several orders of magnitude less memory requirements than the competitors. Only our methods analyze a 5-dimensional tensor with 1 billion observable entries, 10M mode length, and 1K rank, while all others fail (Figure 4(a)).

The binary codes of our methods and several datasets are available at <http://kdmmlab.org/sals>. The rest of this paper is organized as follows. Section II presents preliminaries for tensor factorization and its distributed algorithms. Section III describes our proposed SALS and CDTF methods. Section IV presents the optimization techniques for them on MAPREDUCE. After providing experimental results in Section V, we conclude in Section VI.

II. PRELIMINARIES: TENSOR FACTORIZATION

In this section, we describe the preliminaries of tensor factorization and its distributed algorithms.

A. Tensor and the Notations

Tensors are multi-dimensional arrays that generalize vectors (1-dimensional tensors) and matrices (2-dimensional tensors) to higher dimensions. Like rows and columns in a matrix, an N -dimensional tensor has N modes whose lengths are denoted by I_1 through I_N , respectively. We

denote tensors with variable dimension N by boldface Euler script letters, e.g., \mathcal{X} . Matrices and vectors are denoted by boldface capitals, e.g., \mathbf{A} , and boldface lowercases, e.g., \mathbf{a} , respectively. We denote the entry of a tensor by the symbolic name of the tensor with its indices in subscript. For example, the (i_1, i_2) th entry of \mathbf{A} is denoted by $a_{i_1 i_2}$, and the (i_1, \dots, i_N) th entry of \mathcal{X} is denoted by $x_{i_1 \dots i_N}$. The i_1 th row of \mathbf{A} is denoted by \mathbf{a}_{i_1*} , and the i_2 th column of \mathbf{A} is denoted by \mathbf{a}_{*i_2} . Table III lists the symbols used in this paper.

B. Tensor Factorization

There are several ways to define a tensor factorization. Our definition is based on the PARAFAC decomposition [13], which is one of the most popular decomposition methods, and the nonzero squared loss with L_2 regularization, whose weighted form has been successfully used in many recommender systems [1], [2], [3].

Definition 1 (Tensor Factorization):

Given an N -dimensional tensor $\mathcal{X} (\in \mathbb{R}^{I_1 \times I_2 \times \dots \times I_N})$ with observable entries $\{x_{i_1 \dots i_N} | (i_1, \dots, i_N) \in \Omega\}$, the rank K factorization of \mathcal{X} is to find factor matrices $\{\mathbf{A}^{(n)} \in \mathbb{R}^{I_n \times K} | 1 \leq n \leq N\}$ which minimize the following loss function:

$$L(\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}) = \sum_{(i_1, \dots, i_N) \in \Omega} \left(x_{i_1 \dots i_N} - \sum_{k=1}^K \prod_{n=1}^N a_{i_n k}^{(n)} \right)^2 + \lambda \sum_{n=1}^N \|\mathbf{A}^{(n)}\|_F^2 \quad (1)$$

Note that the loss function depends only on the observable entries. Each factor matrix $\mathbf{A}^{(n)}$ corresponds to the latent feature vectors of the objects that the n th mode of \mathcal{X} represents, and $\sum_{k=1}^K \prod_{n=1}^N a_{i_n k}^{(n)}$ corresponds to the interaction among the features.

C. Distributed Methods for Tensor Factorization

In this section, we explain how widely-used distributed optimization methods are applied to tensor factorization. Their performances are summarized in Table II. Note that only our proposed CDTF and SALS methods, which will be described in Sections III and IV, have no bottleneck in any aspects.

1) *Alternating Least Square (ALS):* Using ALS [3], we update factor matrices one by one while keeping all other matrices fixed. When all other factor matrices are fixed, minimizing (1) is analytically solvable in terms of the updated matrix, which can be updated row by row due to

the independence between rows. The update rule for each row of $\mathbf{A}^{(n)}$ is as follows:

$$\begin{aligned} [a_{i_n 1}^{(n)}, \dots, a_{i_n K}^{(n)}]^T &\leftarrow \arg \min_{[a_{i_n 1}^{(n)}, \dots, a_{i_n K}^{(n)}]^T} L(\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}) \\ &= (\mathbf{B}_{i_n}^{(n)} + \lambda \mathbf{I}_K)^{-1} \mathbf{c}_{i_n}^{(n)} \end{aligned} \quad (2)$$

where the (k_1, k_2) th entry of $\mathbf{B}_{i_n}^{(n)} (\in \mathbb{R}^{K \times K})$ is

$$\sum_{(i_1, \dots, i_N) \in \Omega_{i_n}^{(n)}} \left(\prod_{l \neq n} a_{i_l k_1}^{(l)} \prod_{l \neq n} a_{i_l k_2}^{(l)} \right),$$

the k th entry of $\mathbf{c}_{i_n}^{(n)} (\in \mathbb{R}^K)$ is

$$\sum_{(i_1, \dots, i_N) \in \Omega_{i_n}^{(n)}} \left(x_{i_1 \dots i_N} \prod_{l \neq n} a_{i_l k}^{(l)} \right),$$

and \mathbf{I}_K is the K by K identity matrix. $\Omega_{i_n}^{(n)}$ denotes the subset of Ω whose n th mode's index is i_n . This update rule can be proven as in Theorem 1 in Section III-A since ALS is a special case of SALS. Updating a row, $\mathbf{a}_{i_n}^{(n)}$, for example, using (2) takes $O(|\Omega_{i_n}^{(n)}|K(N+K) + K^3)$, which consists of $O(|\Omega_{i_n}^{(n)}|NK)$ to calculate $\prod_{l \neq n} a_{i_l 1}^{(l)}$ through $\prod_{l \neq n} a_{i_l K}^{(l)}$ for all the entries in $\Omega_{i_n}^{(n)}$, $O(|\Omega_{i_n}^{(n)}|K^2)$ to build $\mathbf{B}_{i_n}^{(n)}$, $O(|\Omega_{i_n}^{(n)}|K)$ to build $\mathbf{c}_{i_n}^{(n)}$, and $O(K^3)$ to invert $\mathbf{B}_{i_n}^{(n)}$. Thus, updating every row of every factor matrix once, which corresponds to a full ALS iteration, takes $O(|\Omega|NK(N+K) + K^3 \sum_{n=1}^N I_n)$.

In distributed environments, updating each factor matrix can be parallelized without affecting the correctness of ALS by distributing the rows of the factor matrix across machines and updating them simultaneously. The parameters updated by each machine are broadcast to all other machines. The number of parameters each machine exchanges with the others is $O(KI_n)$ for each factor matrix $\mathbf{A}^{(n)}$ and $O(K \sum_{n=1}^N I_n)$ per iteration. The memory requirements of ALS, however, cannot be distributed. Since the update rule (2) possibly depends on any entry of any fixed factor matrix, every machine is required to load all the fixed matrices into its memory. This high memory requirements of ALS, $O(K \sum_{n=1}^N I_n)$ memory space per machine, have been noted as a scalability bottleneck even in matrix factorization [4], [5].

2) *Parallelized Stochastic Gradient Descent (PSGD)*: PSGD [9] is a distributed algorithm based on stochastic gradient descent (SGD). In PSGD, the observable entries of \mathcal{X} are randomly divided into M machines which run SGD independently using the assigned entries. The updated parameters are averaged after each iteration. For each observable entry $x_{i_1 \dots i_N}$, $a_{i_n k}^{(n)}$ for all n and k , whose number is NK , are updated at once by the following rule :

$$a_{i_n k}^{(n)} \leftarrow a_{i_n k}^{(n)} - 2\eta \left(\frac{\lambda a_{i_n k}^{(n)}}{|\Omega_{i_n}^{(n)}|} - r_{i_1 \dots i_N} \left(\prod_{l \neq n} a_{i_l k}^{(l)} \right) \right) \quad (3)$$

where $r_{i_1 \dots i_N} = x_{i_1 \dots i_N} - \sum_{s=1}^K \prod_{l=1}^N a_{i_l s}^{(l)}$. It takes $O(NK)$ to calculate $r_{i_1 \dots i_N}$ and $\prod_{l=1}^N a_{i_l k}^{(l)}$ for all k . Once they are calculated, since $\prod_{l \neq n} a_{i_l k}^{(l)}$ can be calculated as $(\prod_{l=1}^N a_{i_l k}^{(l)}) / a_{i_n k}^{(n)}$, calculating (3) takes $O(1)$, and thus updating all the NK parameters takes $O(NK)$. If we assume that \mathcal{X} entries are equally distributed across machines, the computational complexity per iteration is $O(|\Omega|NK/M)$. Averaging parameters can also be distributed, and in the process, $O(K \sum_{n=1}^N I_n)$ parameters are exchanged by each machine. Like ALS, the memory requirements of PSGD cannot be distributed, i.e., all the machines are required to load all the factor matrices into their memory. $O(K \sum_{n=1}^N I_n)$ memory space is required per machine. Moreover, PSGD tends to converge slowly due to the non-identifiability of (1) [4].

3) *Flexible Factorization of Coupled Tensors (Flexi-FaCT)*: FLEXIFACT [10] is another SGD-based algorithm that remedies the high memory requirements and slow convergence of PSGD. FLEXIFACT divides \mathcal{X} into M^N blocks. Each M disjoint blocks that do not share common fibers (rows in a general n th mode) compose a stratum. FLEXIFACT processes \mathcal{X} one stratum at a time in which the M blocks composing a stratum are distributed across machines and processed independently. The update rule is the same as (3), and the computational complexity per iteration is $O(|\Omega|NK/M)$ as in PSGD. However, contrary to PSGD, averaging is unnecessary because a set of parameters updated by each machine are disjoint with those updated by the other machines. In addition, the memory requirements of FLEXIFACT are distributed among the machines. Each machine only needs to load the parameters related to the block it processes, whose number is $(K \sum_{n=1}^N I_n)/M$, into its memory at a time. However, FLEXIFACT suffers from high communication cost. After processing one stratum, each machine sends the updated parameters to the machine which updates them using the next stratum. Each machine exchanges at most $(K \sum_{n=2}^N I_n)/M$ parameters per stratum and $M^{N-2} K \sum_{n=2}^N I_n$ per iteration where M^{N-1} is the number of strata. Thus, the communication cost increases exponentially with the dimension of \mathcal{X} and polynomially with the number of machines.

III. PROPOSED METHODS

In this section, we propose Subset Alternating Least Square (SALS) and Coordinate Descent for Tensor Factorization (CDTF). They are scalable algorithms for tensor factorization, which is essentially an optimization problem whose loss function is (1) and parameters are the entries of factor matrices, $\mathbf{A}^{(1)}$ through $\mathbf{A}^{(N)}$. Figure 1 depicts the difference among CDTF, SALS, and ALS. Unlike ALS, which updates each K columns of factor matrices row by row, SALS updates each C ($1 \leq C \leq K$) columns row by row, and CDTF updates each column entry by entry.




Figure 1: Update rules of CDTF, SALS, and ALS. CDTF updates each column of factor matrices entry by entry, SALS updates each C ($1 \leq C \leq K$) columns row by row, and ALS updates each K columns row by row.

Algorithm 1: Serial version of SALS

Input : \mathcal{X}, K, λ
Output: $\mathbf{A}^{(n)}$ for all n

```

1 initialize  $\mathcal{R}$  and  $\mathbf{A}^{(n)}$  for all  $n$ 
2 for outer iter = 1.. $T_{out}$  do
3   for split iter = 1.. $\lceil \frac{K}{C} \rceil$  do
4     choose  $k_1, \dots, k_C$  (from columns not updated yet)
5     compute  $\hat{\mathcal{R}}$ 
6     for inner iter = 1.. $T_{in}$  do
7       for  $n = 1..N$  do
8         for  $i_n = 1..I_n$  do
9           update  $a_{i_n k_1}^{(n)}, \dots, a_{i_n k_C}^{(n)}$  using (4)
10    update  $\mathcal{R}$  using (5)
```

CDTF can be seen as an extension of CCD++ [5] to higher dimensions. Since SALS contains CDTF ($C = 1$) as well as ALS ($T_{in} = 1, C = K$) as a special case, we focus on SALS then explain additional optimization schemes for CDTF.

A. Update Rule and Update Sequence

Algorithm 1 describes the procedure of SALS. \mathcal{R} denotes the residual tensor where $r_{i_1 \dots i_N} = x_{i_1 \dots i_N} - \sum_{k=1}^K \prod_{n=1}^N a_{i_n k}^{(n)}$. We initialize the entries of $\mathbf{A}^{(1)}$ to zeros and those of all other factor matrices to random values so that the initial value of \mathcal{R} is equal to \mathcal{X} (line 1). In every iteration (line 1), SALS repeats choosing C columns, k_1 through k_C , randomly without replacement (line 1) and updating them while keeping the other columns fixed, which is equivalent to the rank C factorization of $\hat{\mathcal{R}}$ where $\hat{r}_{i_1 \dots i_N} = r_{i_1 \dots i_N} + \sum_{c=1}^C \prod_{n=1}^N a_{i_n k_c}^{(n)}$. Once $\hat{\mathcal{R}}$ is computed (line 1), updating C columns of factor matrices matrix by matrix (line 1) is repeated T_{in} times (line 1). For each factor matrix, since its rows are independent of each other in minimizing (1) when the other factor matrices are fixed, the entries are updated row by row (line 1) as follows:

$$\begin{aligned}
 [a_{i_n k_1}^{(n)}, \dots, a_{i_n k_C}^{(n)}]^T &\leftarrow \arg \min_{[a_{i_n k_1}^{(n)}, \dots, a_{i_n k_C}^{(n)}]^T} L(\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}) \\
 &= (\mathbf{B}_{i_n}^{(n)} + \lambda \mathbf{I}_C)^{-1} \mathbf{c}_{i_n}^{(n)}, \quad (4)
 \end{aligned}$$

where the (c_1, c_2) th entry of $\mathbf{B}_{i_n}^{(n)} (\in \mathbb{R}^{C \times C})$ is

$$\sum_{(i_1, \dots, i_N) \in \Omega_{i_n}^{(n)}} \left(\prod_{l \neq n} a_{i_l k_{c_1}}^{(l)} \prod_{l \neq n} a_{i_l k_{c_2}}^{(l)} \right),$$

the c th entry of $\mathbf{c}_{i_n}^{(n)} (\in \mathbb{R}^C)$ is

$$\sum_{(i_1, \dots, i_N) \in \Omega_{i_n}^{(n)}} \left(\hat{r}_{i_1 \dots i_N} \prod_{l \neq n} a_{i_l k_c}^{(l)} \right),$$

and \mathbf{I}_C is the C by C identity matrix. $\Omega_{i_n}^{(n)}$ denotes the subset of Ω whose n th mode's index is i_n . The proof of this update rule is as follows:

Theorem 1:

$$\arg \min_{[a_{i_n k_1}^{(n)}, \dots, a_{i_n k_C}^{(n)}]^T} L(\mathbf{A}^{(1)}, \dots, \mathbf{A}^{(N)}) = (\mathbf{B}_{i_n}^{(n)} + \lambda \mathbf{I}_C)^{-1} \mathbf{c}_{i_n}^{(n)}$$

Proof:

$$\begin{aligned}
 \frac{\partial L}{\partial a_{i_n k_c}^{(n)}} &= 0, \forall c, 1 \leq c \leq C \\
 \Leftrightarrow \sum_{(i_1, \dots, i_N) \in \Omega_{i_n}^{(n)}} &\left(\left(\sum_{s=1}^C \prod_{l=1}^N a_{i_l k_s}^{(l)} - \hat{r}_{i_1 \dots i_N} \right) \prod_{l \neq n} a_{i_l k_c}^{(l)} \right) + \lambda a_{i_n k_c}^{(n)} \\
 &= 0, \forall c \\
 \Leftrightarrow \sum_{(i_1, \dots, i_N) \in \Omega_{i_n}^{(n)}} &\left(\sum_{s=1}^C \left(a_{i_n k_s}^{(n)} \prod_{l \neq n} a_{i_l k_s}^{(l)} \right) \prod_{l \neq n} a_{i_l k_c}^{(l)} \right) + \lambda a_{i_n k_c}^{(n)} \\
 &= \sum_{(i_1, \dots, i_N) \in \Omega_{i_n}^{(n)}} \left(\hat{r}_{i_1 \dots i_N} \prod_{l \neq n} a_{i_l k_c}^{(l)} \right), \forall c \\
 \Leftrightarrow (\mathbf{B}_{i_n}^{(n)} + \lambda \mathbf{I}_C) [a_{i_n k_1}^{(n)}, \dots, a_{i_n k_C}^{(n)}]^T &= \mathbf{c}_{i_n}^{(n)}
 \end{aligned}$$

Since $\mathbf{B}_{i_n}^{(n)}$ is symmetric, instead of computing its inverse, the Cholesky decomposition can be used. After this rank C factorization, the entries of \mathcal{R} are updated by the following rule (line 1):

$$r_{i_1 \dots i_N} \leftarrow \hat{r}_{i_1 \dots i_N} - \sum_{c=1}^C \prod_{n=1}^N a_{i_n k_c}^{(n)}. \quad (5)$$

CDTF is a special case of SALS where C is set to one. In CDTF, instead of computing $\hat{\mathcal{R}}$ before rank one factorization, the entries of $\hat{\mathcal{R}}$ can be computed while computing (4) and (5). This can result in better performance on a disk-based system like MAPREDUCE by reducing disk I/O operations. Moreover, instead of randomly changing the order by which columns are updated at each iteration, fixing the order speeds up the convergence of CDTF in our experiments.

B. Complexity Analysis

Theorem 2: The computational complexity of Algorithm 1 is $O(T_{out} |\Omega| N T_{in} K (N + C) + T_{out} T_{in} K C^2 \sum_{n=1}^N I_n)$.

Proof:

Computing $\hat{\mathcal{R}}$ (line 1) and updating \mathcal{R} (line 1) take $O(|\Omega| N C)$. Updating C parameters (line 1) takes $O(|\Omega_{i_n}^{(n)}| C (C + N) + C^3)$, which consists of $O(|\Omega_{i_n}^{(n)}| N C)$

Algorithm 2: Distributed version of SALS

Input : \mathcal{X} , K , λ , mS_n for all m and n

Output: $\mathbf{A}^{(n)}$ for all n

```

1 distribute the  $m\Omega$  entries of  $\mathcal{X}$  to each machine  $m$ 
2 Parallel (P): initialize the  $m\Omega$  entries of  $\mathcal{R}$ 
3 P: initialize  $\mathbf{A}^{(n)}$  for all  $n$ 
4 for outer iter = 1.. $T_{out}$  do
5   for split iter = 1.. $\lceil \frac{K}{C} \rceil$  do
6     choose  $k_1, \dots, k_C$  (from columns not updated yet)
7     P: compute  $m\Omega$  entries of  $\hat{\mathcal{R}}$ 
8     for inner iter = 1.. $T_{in}$  do
9       for  $n = 1..N$  do
10        P: update  $\{a_{i_n k_c}^{(n)} | i_n \in mS_n, 1 \leq c \leq C\}$ 
11        using (4)
12        P: broadcast  $\{a_{i_n k_c}^{(n)} | i_n \in mS_n, 1 \leq c \leq C\}$ 
13      P: update the  $m\Omega$  entries of  $\mathcal{R}$  using (5)

```

to calculate $\prod_{l \neq n} a_{i_l k_1}^{(l)}$ through $\prod_{l \neq n} a_{i_l k_C}^{(l)}$ for all the entries in $|\Omega_{i_n}^{(n)}|$, $O(|\Omega_{i_n}^{(n)}|C^2)$ to build $\mathbf{B}_{i_n}^{(n)}$, $O(|\Omega_{i_n}^{(n)}|C)$ to build $\mathbf{c}_{i_n}^{(n)}$, and $O(C^3)$ to invert $\mathbf{B}_{i_n}^{(n)}$. Thus, updating all the entries in C columns (lines 1 through 1) takes $O(|\Omega|C(C+N) + I_n C^3)$, and the rank C factorization (lines 1 through 1) takes $O(|\Omega|NC(N+C) + C^3 \sum_{n=1}^N I_n)$. As a result, an outer iteration, which repeats the rank C factorization $T_{in}K/C$ times and both $\hat{\mathcal{R}}$ computation and \mathcal{R} update K/C times, takes $O(|\Omega|NT_{in}K(N+C) + T_{in}KC^2 \sum_{n=1}^N I_n) + O(|\Omega|NK)$, where the second term is dominated. ■

Theorem 3: The memory requirement of Algorithm 1 is $O(C \sum_{n=1}^N I_n)$.

Proof:

Since $\hat{\mathcal{R}}$ computation (line 1), rank C factorization (lines 1 through 1), and \mathcal{R} update (line 1) all depend only on the C columns of the factor matrices, the number of whose entries is $C \sum_{n=1}^N I_n$, the other $(K-C)$ columns do not need to be loaded into memory. Thus, the columns of the factor matrices can be loaded by turns depending on (k_1, \dots, k_C) values. Moreover, updating C columns (lines 1 through 1) can be processed by streaming the entries of $\hat{\mathcal{R}}$ from disk and processing them one by one instead of loading them all at once because the entries of $\mathbf{B}_{i_n}^{(n)}$ and $\mathbf{c}_{i_n}^{(n)}$ in (4) are the sum of the values calculated independently from each $\hat{\mathcal{R}}$ entry. Likewise, $\hat{\mathcal{R}}$ computation and \mathcal{R} update can also be processed by streaming \mathcal{R} and $\hat{\mathcal{R}}$, respectively. ■

C. Parallelization in Distributed Environments

In this section, we describe how to parallelize SALS in distributed environments such as MAPREDUCE where machines do not share memory. Algorithm 2 depicts the distributed version of SALS.

Since update rule (4) for each row (C parameters) of a factor matrix does not depend on the other rows in the matrix, rows in a factor matrix can be distributed across machines




Figure 2: Work and data distribution of SALS in distributed environments when an input tensor is a three-dimensional tensor and the number of machines is four. We assume that the rows of the factor matrices are assigned to the machines sequentially. The colored region of $\bar{\mathbf{A}}^{(n)}$ (the transpose of $\mathbf{A}^{(n)}$) in each sub-figure corresponds to the parameters updated by each machine, resp., and that of \mathcal{X} corresponds to the data distributed to each machine.

and updated simultaneously without affecting the correctness of SALS. Each machine m updates mS_n rows of $\mathbf{A}^{(n)}$ (line 2), and for this, the $m\Omega = \bigcup_{n=1}^N \left(\bigcup_{i_n \in mS_n} \Omega_{i_n}^{(n)} \right)$ entries of \mathcal{X} are distributed to machine m in the first stage of the algorithm (line 2). Figure 2 shows an example of work and data distribution in SALS.

After the update, parameters updated by each machine are broadcast to all other machines (line 2). Each machine m broadcasts $C|mS_n|$ parameters and receives $C(I_n - |mS_n|)$ parameters from the other machines after each update. The total number of parameters each machine exchanges with the other machines is $KT_{in} \sum_{n=1}^N I_n$ per outer iteration.

The running time of parallel steps in Algorithm 2 depends on the longest running time among all machines. Specifically, the running time of lines 2, 2, and 2 is proportional to $\max_m |m\Omega^{(n)}|$ where $m\Omega^{(n)} = \bigcup_{i_n \in mS_n} \Omega_{i_n}^{(n)}$, and that of line 2 is proportional to $\max_m |mS_n|$. Therefore, it is necessary to assign the rows of the factor matrices to the machines (i.e., to decide mS_n) so that $|m\Omega^{(n)}|$ and $|mS_n|$ are even among all the machines. The greedy assignment algorithm described in Algorithm 3 aims to minimize $\max_m |m\Omega^{(n)}|$ under the condition that $\max_m |mS_n|$ is minimized (i.e., $|mS_n| = I_n/M$ for all n where M is the number of machines). For each factor matrix $\mathbf{A}^{(n)}$, we sort its rows in the decreasing order of $|\Omega_{i_n}^{(n)}|$ and assign the rows one by one to the machine m which satisfies $|mS_n| < \lceil I_n/M \rceil$ and has the smallest $|m\Omega^{(n)}|$ currently. The effects of this greedy row assignment on actual running times are analyzed in Section V-E.

IV. OPTIMIZATION ON MAPREDUCE

In this section, we describe two optimization techniques used to implement SALS and CDTF on MAPREDUCE, which is one of the most widely-used distributed platforms.

A. Local Disk Caching

Typical MAPREDUCE implementations of SALS and CDTF without local disk caching run each parallel step as a separate MAPREDUCE job. Algorithms 4 and 5 describe the MAPREDUCE implementation of parameter update (update

Algorithm 3: Greedy row assignment in SALS

Input : \mathcal{X}, M
Output: $_m S_n$ for all m and n

```
1 initialize  $|_m \Omega|$  to 0 for all  $m$ 
2 for  $n = 1..N$  do
3   initialize  $_m S_n$  to  $\emptyset$  for all  $m$ 
4   initialize  $|_m \Omega^{(n)}|$  to 0 for all  $m$ 
5   calculate  $|\Omega_{i_n}^{(n)}|$  for all  $i_n$ 
6   foreach  $i_n$  (in decreasing order of  $|\Omega_{i_n}^{(n)}|$ ) do
7     find  $m$  with  $|_m S_n| < \lceil \frac{I_n}{M} \rceil$  and the smallest  $|_m \Omega^{(n)}|$ 
8     (in case of a tie, choose the machine with smaller  $|_m S_n|$ ,
9     and if still a tie, choose the one with smaller  $|_m \Omega|$ )
10    add  $i_n$  to  $_m S_n$ 
11    add  $|\Omega_{i_n}^{(n)}|$  to  $|_m \Omega^{(n)}|$  and  $|_m \Omega|$ 
```

Algorithm 4: Parameter update in SALS without local disk caching

Given : n, k_c for all c , $_m S_n$ for all m , $\mathbf{a}_{*k_c}^{(l)}$ for all l and c
Input : $\hat{\mathcal{R}}$
Output: updated $\mathbf{a}_{*k_c}^{(n)}$ for all c

```
1 Map(Key k, Value v)
2 begin
3    $((i_1, \dots, i_N), \hat{r}_{i_1 \dots i_N}) \leftarrow v$ 
4   find  $m$  where  $i_n \in _m S_n$ 
5   emit  $\langle (m, i_n), ((i_1, \dots, i_N), \hat{r}_{i_1 \dots i_N}) \rangle$ 
6 end
7 Partitioner(Key k, Value v)
8 begin
9    $(m, i_n) \leftarrow k$ 
10  assign  $\langle k, v \rangle$  to machine  $m$ 
11 end
12 Reduce(Key k, Value v[1..r])
13 begin
14    $(m, i_n) \leftarrow k$ 
15    $\Omega_{i_n}^{(n)}$  entries of  $\hat{\mathcal{R}} \leftarrow v$ 
16   update and emit  $\mathbf{a}_{i_n k_c}^{(n)}$  for all  $c$ 
17 end
```

of $\mathbf{a}_{*k_c}^{(n)}$ for all c) and \mathcal{R} update, respectively. $\hat{\mathcal{R}}$ computation can be implemented by the similar way with \mathcal{R} update. In this implementation, broadcasting updated parameters is unnecessary because reducers terminate after updating their assigned parameters. Instead, the updated parameters are saved in the distributed file system and are read at the next step (a separate job). Since SALS repeats both \mathcal{R} update and $\hat{\mathcal{R}}$ computation K/C times and parameter update $KT_{in}N/C$ times at every outer iteration, this implementation repeats distributing \mathcal{R} or $\hat{\mathcal{R}}$ across machines (the mapper stage of Algorithms 4 and 5) $T_{out}K(T_{in}N + 2)/C$ times, which is inefficient.

Our implementation reduces this inefficiency by caching data to local disk once they are distributed. In the SALS implementation with local disk caching, \mathcal{X} entries are distributed across machines and cached in the local disk during the map and reduce stages (Algorithm 6); and the rest part

Algorithm 5: \mathcal{R} update in SALS without local disk caching

Given : k_c for all c , $_m S_1$ for all m , $\mathbf{a}_{*k_c}^{(l)}$ for all l and c
Input : $\hat{\mathcal{R}}$
Output: updated \mathcal{R}

```
1 Map(Key k, Value v)
2 begin
3    $((i_1, \dots, i_N), \hat{r}_{i_1 \dots i_N}) \leftarrow v$ 
4   find  $m$  where  $i_1 \in _m S_1$ 
5   emit  $\langle m, ((i_1, \dots, i_N), \hat{r}_{i_1 \dots i_N}) \rangle$ 
6 end
7 Partitioner(Key k, Value v)
8 begin
9    $m \leftarrow k$ 
10  assign  $\langle k, v \rangle$  to machine  $m$ 
11 end
12 Reduce(Key k, Value v[1..r])
13 begin
14   foreach  $((i_1, \dots, i_N), \hat{r}_{i_1 \dots i_N}) \in v[1..r]$  do
15     update  $r_{i_1 \dots i_N}$ 
16     emit  $((i_1, \dots, i_N), r_{i_1 \dots i_N})$ 
17 end
```

Algorithm 6: Data distribution in SALS with local disk caching

Input : $\mathcal{X}, _m S_n$ for all m and n
Output: $_m \Omega^{(n)}$ entries of $\mathcal{R}(=\mathcal{X})$ for all m and n

```
1 Map(Key k, Value v);
2 begin
3    $((i_1, \dots, i_N), x_{i_1 \dots i_N}) \leftarrow v$ 
4   for  $n = 1, \dots, N$  do
5     find  $m$  where  $i_n \in _m S_n$ 
6     emit  $\langle (m, n), ((i_1, \dots, i_N), x_{i_1 \dots i_N}) \rangle$ 
7   end
8 Partitioner(Key k, Value v);
9 begin
10   $(m, n) \leftarrow k$ 
11  assign  $\langle k, v \rangle$  to machine  $m$ 
12 end
13 Reduce(Key k, Value v[1..r]);
14 begin
15    $(m, n) \leftarrow k$ 
16   open a file on the local disk to cache  $_m \Omega^{(n)}$  entries of  $\mathcal{R}$ 
17   foreach  $((i_1, \dots, i_N), x_{i_1 \dots i_N}) \in v[1..r]$  do
18     write  $((i_1, \dots, i_N), x_{i_1 \dots i_N})$  to the file
19 end
```

of SALS runs in the close stage (cleanup stage in Hadoop) using cached data. Our implementation streams the cached data from the local disk instead of distributing entire \mathcal{R} or $\hat{\mathcal{R}}$ from the distributed file system when updating factor matrices. For example, $_m \Omega^{(n)}$ entries of $\hat{\mathcal{R}}$ are streamed from the local disk when the columns of $\mathbf{A}^{(n)}$ are updated. The effect of this local disk caching on the actual running time is analyzed in Section V-E.

Algorithm 7: Parameter broadcast in SALS

Input : $m\mathbf{a}_{*k_c}^{(n)}$ for all c (parameters to broadcast)
Output: $\mathbf{a}_{*k_c}^{(n)}$ for all c (parameters received from others)
1 **begin**
2 create a data file mA on the distributed file system (DFS)
3 write $m\mathbf{a}_{*k_c}^{(n)}$ on the data file
4 create a dummy file mD on DFS
5 **while** not all data files were read **do**
6 get the list of dummy files from DFS
7 **foreach** $m'D$ in the list **do**
8 **if** $m'A$ were not read **then**
9 read $m'\mathbf{a}_{*k_c}^{(n)}$ from $m'A$
10 **end**

B. Direct Communication

In MAPREDUCE, it is generally assumed that reducers run independently and do not communicate directly with each other. However, we adapt the direct communication method using the distributed file system in [10] to broadcast parameters among reducers efficiently. The implementation of parameter broadcast in SALS (i.e., broadcast of $\mathbf{a}_{*k_c}^{(n)}$ for all c) based on this method is described in Algorithm 7 where $m\mathbf{a}_{*k_c}^{(n)}$ denotes $\{a_{i_n k_c}^{(n)} | i_n \in mS_n\}$.

V. EXPERIMENTS

To evaluate SALS and CDTF, we design and conduct experiments to answer the following questions:

- **Q1: Data scalability (Section V-B).** How do SALS, CDTF, and other methods scale with regard to the following properties of an input tensor: dimension, the number of observations, mode length, and rank?
- **Q2: Machine scalability (Section V-C).** How do SALS, CDTF, and other methods scale with regard to the number of machines?
- **Q3: Convergence (Section V-D).** How quickly and accurately do SALS, CDTF, and other methods factorize real-world tensors?
- **Q4: Optimization (Section V-E).** How much do the local disk caching and the greedy row assignment improve the speed of SALS and CDTF? Can these optimization techniques be applied to other methods?
- **Q5: Effects of T_{in} (Section V-F)** How do different numbers of inner iterations (T_{in}) affect the convergence of CDTF?
- **Q6: Effects of C (Section V-G)** How do different numbers of columns updated at a time (C) affect the convergence of SALS?

Other methods include ALS, FLEXIFACT, and PSGD, which are explained in Section II-C. All experiments are focused on the distributed version of each method, which is the most suitable to achieve our purpose of handling large-scale data.

Table IV: Summary of real-world datasets.

	Movielens ₄	Netflix ₃	Yahoo-music ₄
N	4	3	4
I_1	715,670	2,649,429	1,000,990
I_2	65,133	17,770	624,961
I_3	169	74	133
I_4	24	-	24
$ \Omega $	93,012,740	99,072,112	252,800,275
$ \Omega _{test}$	6,987,800	1,408,395	4,003,960
K	20	40	80
λ	0.01	0.02	1.0
η_0	0.01	0.01	10^{-5} (FLEXIFACT) 10^{-4} (PSGD)

Table V: Scale of synthetic datasets. B: billion, M: million, K: thousand. The length of every mode is equal to I .

	S1	S2 (default)	S3	S4
N	2	3	4	5
I	300K	1M	3M	10M
$ \Omega $	30M	100M	300M	1B
K	30	100	300	1K

A. Experimental Settings

1) *Cluster:* We run experiments on a 40-node Hadoop cluster. Each node has an Intel Xeon E5620 2.4GHz CPU. The maximum heap memory size per reducer is set to 8GB.

2) *Data:* We use both real-world and synthetic datasets most of which are available at <http://kdmlab.org/sals>. The real-world tensor data used in our experiments are summarized in Table IV with the following details:

- **Movielens₄¹:** Movie rating data from MovieLens, an online movie recommender service. We convert them into a four-dimensional tensor where the third mode and the fourth mode correspond to (year, month) and hour-of-day when the movie is rated, respectively. The rates range from 1 to 5.
- **Netflix₃²:** Movie rating data used in Netflix prize. We regard them as a three-dimensional tensor where the third mode corresponds to (year, month) when the movie is rated. The rates range from 1 to 5.
- **Yahoo-music₄³:** Music rating data used in KDD CUP 2011. We convert them into a four-dimensional tensor by the same way in Movielens₄. Since exact year and month are not provided, we use the values obtained by dividing the provided data (the number of days passed from an unrevealed date) by 30. The rates range from 0 to 100.

For reproducibility, we use the original training/test split offered by the data providers. Synthetic tensors are created by the procedure used in [14] to create Jumbo dataset. The scales of the synthetic datasets used are summarized in Table V.

3) *Implementation and Parameter Settings:* All the methods in Table II are implemented in Java with Hadoop 1.0.3.

¹<http://grouplens.org/datasets/movielens>

²<http://www.netflixprize.com>

³<http://webscope.sandbox.yahoo.com/catalog.php?datatype=c>




Figure 3: Scalability with regard to each aspect of data. o.o.m. : out of memory. Only SALS and CDTF scale with all the aspects.

The local disk caching, the direct communication, and the greedy row assignment are applied to all the methods if possible. All our implementations use weighted- λ -regularization [3]. For SALS and CDTF, T_{in} is set to 1, and C is set to 10, unless otherwise stated. The learning rate of FLEXIFACT and PSGD at t th iteration is set to $2\eta_0/(1+t)$ following the open-sourced FLEXIFACT implementation (<http://alexbeutel.com/flexifact/>). The number of reducers is set to 5 for FLEXIFACT, 20 for PSGD, and 40 for the other methods, each of which leads to the best performance on the machine scalability test in Section V-C, unless otherwise stated.

B. Data Scalability

1) *Scalability with Each Factor (Figure 3):* We measure the scalability of CDTF, SALS, and the competitors with regard to the dimension, number of observations, mode length, and rank of an input tensor. When measuring the scalability with regard to a factor, the factor is scaled up from S1 to S4 while all other factors are fixed at S2 as summarized in Table V. As seen in Figure 3(a), FLEXIFACT does not scale with dimension because of its communication cost, which increases exponentially with dimension. ALS and PSGD are not scalable with mode length and rank due to their high memory requirements as Figures 3(c) and 3(d) show. They require up to 11.2GB, which is $48\times$ of 234MB that CDTF requires and $10\times$ of 1,147MB that SALS requires. Moreover, the running time of ALS increases rapidly with rank owing to its cubically increasing computational cost. Only SALS and CDTF are scalable with all the factors as summarized in Table I. Their running times increase linearly with all the factors except dimension, with which they increase slightly faster due to the quadratically increasing computational cost.

2) *Overall Scalability (Figure 4(a)):* We measure the scalability of the methods by scaling up all the factors together from S1 to S4. The scalability of FLEXIFACT with five machines, ALS, and PSGD is limited owing to their high memory requirements. ALS and PSGD require almost 186GB to handle S4, which is $493\times$ of 387MB that CDTF requires and $100\times$ of 1,912MB that SALS requires. FLEXIFACT with 40 machines does not scale over S2 due to its rapidly increasing communication cost. Only SALS and CDTF scale up to S4, and there is a trade-off between them:




Figure 4: (a) Overall scalability. o.o.m. : out of memory, o.o.t. : out of time (takes more than a week). Only SALS and CDTF scale up to the largest scale S4. (b) Machine scalability. Computations of SALS and CDTF are efficiently distributed across machines.

SALS runs faster, and CDTF is more memory-efficient.

C. Machine Scalability (Figure 4(b))

We measure the speed-ups (T_5/T_M where T_M is the running time with M reducers) of the methods on the S2 scale dataset by increasing the number of reducers. The speed-ups of CDTF, SALS, and ALS increase linearly at the beginning and then flatten out slowly owing to their fixed communication cost which does not depend on the number of reducers. The speed-up of PSGD flattens out fast, and PSGD even slightly slows down at 40 reducers because of increased overhead. FLEXIFACT slows down as the number of reducers increases because of its rapidly increasing communication cost.

D. Convergence (Figure 5)

We compare how quickly and accurately each method factorizes real-world tensors. Accuracies are calculated at each iteration by root mean square error (RMSE) on a held-out test set, which is a measure commonly used by recommender systems. Table IV describes K , λ , and η_0 values used for each dataset. They are determined by cross validation. Owing to the non-convexity of (1), each algorithm may converge to local minima with different accuracies. In all datasets (results on the Movielens₄ dataset are omitted for space reasons), SALS is comparable with ALS, which converges the fastest to the best solution, and CDTF follows them. PSGD converges the slowest to the worst solution due to the non-identifiability of (1) [4].




Figure 5: Convergence speed on real-world datasets. SALS is comparable with ALS, which converges fastest to the best solution, and CDTF follows them.




Figure 6: Effects of optimization techniques on running times. NC: no caching, LC: local disk caching, SEQ: sequential row assignment⁴, RAN: random row assignment, GRE: greedy row assignment. Our proposed optimization techniques (LC+GRE) significantly accelerate CDTF, SALS, and also their competitors.

E. Optimization (Figure 6)

We measure how our proposed optimization techniques, the local disk caching and the greedy row assignment, affect the running time of CDTF, SALS, and the competitors on real-world datasets. The direct communication method explained in Section IV-B is applied to all the implementations if necessary. The local disk caching speeds up CDTF up to $65.7\times$, SALS up to $15.5\times$, and the competitors up to $4.8\times$. The speed-ups of SALS and CDTF are the most significant because of the highly iterative nature of SALS and CDTF. Additionally, the greedy row assignment speeds up CDTF up to $1.5\times$; SALS up to $1.3\times$; and the competitors up to $1.2\times$ compared with the second best one. It is not applicable to PSGD, which does not distribute parameters row by row.

F. Effects of Inner Iterations (Figure 7)

We compare the convergence properties of CDTF with different T_{in} values. As T_{in} increases, CDTF tends to converge more stably to better solutions (with lower test RMSE). However, there is an exception, $T_{in} = 1$ in the Netflix₃ dataset, which converges to the best solution.

G. Effects of the Number of Columns Updated at a Time (Figure 8)

We compare the convergence properties of SALS with different C values. T_{in} is fixed to one in all cases. As C increases, SALS tends to converge faster to better solutions

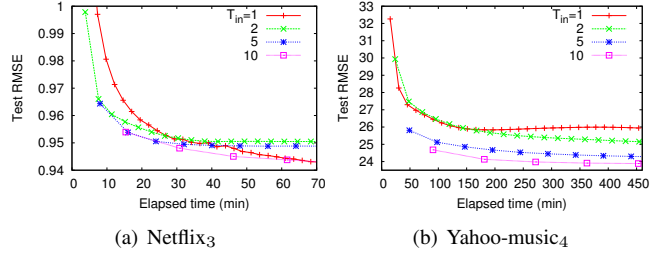


Figure 7: Effects of inner iterations (T_{in}) on the convergence of CDTF. CDTF tends to converge stably to better solutions as T_{in} increases.




Figure 8: Effects of the number of columns updated at a time (C) on the convergence of SALS. SALS tends to converge faster to better solutions as C increases. However, its convergence speed decreases at C above 20.




Figure 9: Effects of the number of columns updated at a time (C) on the running time of SALS. Running time per iteration decreases until $C = 20$, then starts to increase.

(with lower test RMSE) although it requires more memory as explained in Theorem 3. With C above 20, however, convergence speed starts to decrease. This trend is related to the running time per iteration, which shows the same trend as seen in Figure 9. As C increases, the amount of disk I/O declines since it depends on the number of times that the entries of \mathcal{R} or $\hat{\mathcal{R}}$ are streamed from disk, which is inversely proportional to C . Conversely, computational cost increases quadratically with regard to C . At low C values, the decrease in the amount of disk I/O is greater and leads to a downward trend of running time per iteration. The opposite happens at high C values.

VI. CONCLUSION

In this paper, we propose SALS and CDTF, distributed algorithms for high-dimensional large-scale tensor factorization. They are scalable with all aspects of data (i.e.,

⁴ $mS_n = \{i_n \in \mathbb{N} | \frac{I_n \times (m-1)}{M} < i_n \leq \frac{I_n \times m}{M}\}$

dimension, the number of observable entries, mode length, and rank) and show a trade-off: SALS has an advantage in terms of convergence speed, and CDTF has one in terms of memory usage. The local disk caching and the greedy row assignment, two proposed optimization schemes, significantly accelerate not only SALS and CDTF but also their competitors.

ACKNOWLEDGMENTS

This work was supported by AFOSR/AOARD under the Grant No. FA2386-14-1-4036, and by the National Research Foundation of Korea (NRF) Grant funded by the Korean Government (MSIP) (No. 2013R1A1A1064409)

REFERENCES

- [1] P.-L. Chen, C.-T. Tsai, Y.-N. Chen, K.-C. Chou, C.-L. Li, C.-H. Tsai, K.-W. Wu, Y.-C. Chou, C.-Y. Li, W.-S. Lin, *et al.*, “A linear ensemble of individual and blended models for music rating prediction,” *KDDCup 2011 Workshop*, 2011.
- [2] Y. Koren, R. Bell, and C. Volinsky, “Matrix factorization techniques for recommender systems,” *Computer*, vol. 42, no. 8, pp. 30–37, 2009.
- [3] Y. Zhou, D. Wilkinson, R. Schreiber, and R. Pan, “Large-scale parallel collaborative filtering for the netflix prize,” in *AAIM*, pp. 337–348, 2008.
- [4] R. Gemulla, E. Nijkamp, P. Haas, and Y. Sismanis, “Large-scale matrix factorization with distributed stochastic gradient descent,” in *KDD*, 2011.
- [5] C.-J. H. S. S. Yu, Hsiang-Fu and I. S. Dhillon., “Parallel matrix factorization for recommender systems,” *Knowl. Inf. Syst.*, pp. 1–27, 2013.
- [6] A. Karatzoglou, X. Amatriain, L. Baltrunas, and N. Oliver, “Multiverse recommendation: N-dimensional tensor factorization for context-aware collaborative filtering,” in *RecSys*, 2010.
- [7] A. Nanopoulos, D. Rafailidis, P. Symeonidis, and Y. Manolopoulos, “Musicbox: Personalized music recommendation based on cubic analysis of social tags,” *IEEE TASLP*, vol. 18, no. 2, pp. 407–412, 2010.
- [8] V. W. Zheng, B. Cao, Y. Zheng, X. Xie, and Q. Yang, “Collaborative filtering meets mobile recommendation: A user-centered approach,” in *AAAI*, 2010.
- [9] R. McDonald, K. Hall, and G. Mann, “Distributed training strategies for the structured perceptron,” in *HLT-NAACL*, 2010.
- [10] A. Beutel, A. Kumar, E. E. Papalexakis, P. P. Talukdar, C. Faloutsos, and E. P. Xing, “Flexifact: Scalable flexible factorization of coupled tensors on hadoop,” in *SDM*, 2014.
- [11] D. M. Dunlavy, T. G. Kolda, and E. Acar, “Temporal link prediction using matrix and tensor factorizations,” *ACM TKDD*, vol. 5, no. 2, p. 10, 2011.
- [12] J.-T. Sun, H.-J. Zeng, H. Liu, Y. Lu, and Z. Chen, “Cubesvd: a novel approach to personalized web search,” in *WWW*, 2005.
- [13] T. Kolda and B. Bader, “Tensor decompositions and applications,” *SIAM review*, vol. 51, no. 3, 2009.
- [14] F. Niu, B. Recht, C. Ré, and S. J. Wright, “Hogwild!: A lock-free approach to parallelizing stochastic gradient descent,” *NIPS*, 2011.